

10/758,486

IBM/265

In the Claims

1. (original) A method for generating a statistic for a pattern matching predicate on an attribute of a relation, to be used in optimizing execution of a query directed to one or more attributes of said relation, comprising

providing a data storage structure storing character statistics on said attribute, comprising a first structure storing, for each of a plurality of character positions, frequently occurring characters in that character position, and statistics for each frequently occurring character,

retrieving said statistics in response to said pattern matching predicate based upon the character positions of characters in said pattern matching predicate, and

generating said statistic based upon said retrieved character statistics.

2. (original) The method of claim 1 wherein said data storage structure further stores a count of a number of occurrences of a frequently occurring character.

3. (original) The method of claim 1 wherein said data storage structure further comprises a second structure storing frequently occurring characters that are subsequent to the frequently occurring characters stored in said first structure.

4. (original) The method of claim 3 wherein said second data storage structure further stores a probability of occurrence of a frequently occurring subsequent character.

10/758.486

IBM/265

5. (original) The method of claim 1 wherein generating said statistic based upon said retrieved character statistics comprises estimating a statistic for a desired character and desired character position, in the event said statistic is not stored in said first structure.

6. (original) The method of claim 5 wherein said estimating comprises accumulating statistics for said desired character position, for characters other than said desired character, and calculating average frequency of occurrence of characters that do not have statistics stored in said first structure.

7. (original) The method of claim 1 wherein said first structure is a table.

8. (original) The method of claim 1 wherein said first structure comprises a linked list.

9. (original) The method of claim 3 wherein said second structure comprises a table.

10. (original) The method of claim 3 wherein said second structure comprises a linked list.

11. (original) A computer system implementing a relational database system and generating a statistic for a pattern matching predicate on an attribute of a relation of said relational database, to be used in optimizing execution of a query directed to one or more attributes of said relation, comprising

10/758,486

IBM/265

storage for said relational database, including a relation having a plurality of tuples including values for a plurality of attributes, and a data storage structure storing character statistics on an attribute, comprising a first structure storing, for each of a plurality of character positions, frequently occurring characters in that character position, and statistics for each frequently occurring character, and

computing circuitry performing query optimization and query execution upon said relational database, said query optimization including generating said statistic for an attribute of said relation by retrieving said statistics in response to said pattern matching predicate based upon the character positions of characters in said pattern matching predicate, and generating said statistic based upon said retrieved character statistics.

12. (original) The computer system of claim 11 wherein said data storage structure further stores a count of a number of occurrences of a frequently occurring character.

13. (original) The computer system of claim 11 wherein said data storage structure further comprises a second structure storing frequently occurring characters that are subsequent to the frequently occurring characters stored in said first structure.

14. (original) The computer system of claim 13 wherein said second data storage structure further stores a probability of occurrence of a frequently occurring subsequent character.

10/758,486

IBM/265

15. (original) The computer system of claim 11 wherein said computing circuitry generates said statistic based upon said retrieved character statistics by estimating a statistic for a desired character and desired character position, in the event said statistic is not stored in said first structure.

16. (original) The computer system of claim 15 wherein said computer circuitry performs said estimating by accumulating statistics for said desired character position, for characters other than said desired character, and calculating average frequency of occurrence of characters that do not have statistics stored in said first structure.

17. (original) The computer system of claim 11 wherein said first structure is a table.

18. (original) The computer system of claim 11 wherein said first structure comprises a linked list.

19. (original) The computer system of claim 13 wherein said second structure comprises a table.

20. (original) The computer system of claim 13 wherein said second structure comprises a linked list.

10/758,486

IBM/265

21. (original) A program product for implementing a relational database system and generating a statistic for a pattern matching predicate on an attribute of a relation of said relational database, to be used in optimizing execution of a query directed to one or more attributes of said relation, comprising

a relational database, including a relation having a plurality of tuples including values for a plurality of attributes,

a data storage structure storing character statistics on an attribute, comprising a first structure storing, for each of a plurality of character positions, frequently occurring characters in that character position, and statistics for each frequently occurring character, and

relational database software performing query optimization and query execution upon said relational database, said query optimization including generating a statistic for an attribute of said relation by retrieving said statistics in response to said pattern matching predicate based upon the character positions of characters in said pattern matching predicate, and generating said statistic based upon said retrieved character statistics, and

a signal bearing media holding said relational database and relational database software.

22. (original) The program product of claim 21 wherein the signal bearing media comprises transmission media.

10/758,486

IBM/265

23. (original) The program product of claim 21 wherein the signal bearing media comprises recordable media.